


(12) UK Patent Application (19) GB (11) 2 350 211 (13) A

(43) Date of A Publication 22.11.2000

(21) Application No 9900884.9	(61) INT CL ⁷ G06F 17/30, H04L 29/10
(22) Date of Filing 14.01.1999	(52) UK CL (Edition R) G4A AUDB
(71) Applicant(s) Paul William Kuczora 69 St Mildred's Road, WESTGATE-ON-SEA, Kent, CT8 8RL, United Kingdom	(56) Documents Cited EP 0748095 A2 WO 98/48546 A1 PC Magazine: The 1997 Utility Guide (Internet Filtering Utilities)(Editor's Choice) http://www.zdnet.co Comms of Assocn for Computing Machinery, 1996, Vol39 No10, pp.87-93 esnick P "PICS:Internet Access Contr
(72) Inventor(s) Paul William Kuczora	(58) Field of Search UK CL (Edition R) G4A AUDB INT CL ⁷ G06F 17/30, H04L 29/10 Online: EPODOC, JAPIO, WPI / EPOQUE; GOOGLE / Internet
(74) Agent and/or Address for Service Paul William Kuczora 69 St Mildred's Road, WESTGATE-ON-SEA, Kent, CT8 8RL, United Kingdom	

(54) Abstract Title
INTERNET BROWSER SOFTWARE LOCK

(57) A computer software system which combines the function of an Internet portal, providing a structured and categorised access point to the Internet, with that of restricting Internet access for certain end-users on behalf of someone acting in a supervisory role. A preferred implementation of the software operates at the user interface and operating system level, conferring the ability to monitor and control other software systems running on the client computer. This approach allows the software to also resist attempts at tampering and to monitor patterns of use of other software on the system. All system parameters can be modified by the person acting in a supervisory role (such as a parent or office manager) via password-protected access to the system settings and data files.



GB 2 350 211 A

FIGURE 1 - GENERALISED VIEW OF TYPICAL INTERNET CONNECTION TOPOLOGIES





FIGURE 2 - ABSTRACTION OF THE COMMUNICATION LEVELS USED DURING INTERNET ACCESS



INTERNET BROWSER LOCK SOFTWARE

This invention relates to a computer software system which acts to restrict Internet access to a list of prescribed sites, while at the same time providing the user with a structured view of the information sources which are available.

A number of Internet filtering software packages already exist, but they suffer from the disadvantage that they are predominantly based on a "black list" approach where either each "unacceptable" Internet site has to be explicitly listed in order for access to be blocked, or where Internet content is modified by simply excising any words from a given Web page if they happen to be on the proscribed list. This approach requires that the location of all unacceptable Internet sites is continually monitored, but this has proved to constitute a moving target that will always outrun attempts to encompass it. In addition, a "black list" requires that a database of the locations or wording of the very worst content on the Internet is maintained on every user's computer. Such data is obviously encoded, but instructions are freely available regarding how to circumvent the encryption methods which are employed by some of these systems.

At the same time, the need to combat "information overload", by trying to structure the huge and ever-increasing amount of information that is available on the Internet, has resulted in the development of Internet "portals". These are key Internet sites which attempt to act as entry points to the Internet for large numbers of users. Such sites are, however, having problems in satisfying user requirements, as their monolithic and remote (in networking terms) nature makes it difficult to offer truly personalised individual access. The ability to block access to any Internet content which is deemed unacceptable is also impossible to achieve at the portal level.

An object of this invention is to provide a computer software mechanism which restricts Internet access to a prescribed list of sites, as defined by some supervisor (such as a parent or office manager), while simultaneously presenting the end user with a highly structured, categorised view of the sites that are available to be accessed.

Accordingly, this invention provides a mechanism whereby the same software system is able to perform both the function of providing a personalised, user-specific Internet portal, while simultaneously barring access to any Internet sites which are not listed on that portal. The software is to be installed on a target computer by someone acting in a supervisory role, who is provided with password-protected access to all parameters used by the software.

In addition to combining the functions of Web-filtering and providing a structured, user-specific Internet portal, the software is intended to employ an entirely different method of implementing the "access locking" function, in comparison to the methods used by existing software systems. It achieves this by monitoring and controlling computer activity at the user interface level, rather than at the network protocol level, and does this without requiring any modifications to any elements of the underlying computer Operating System. Such an approach is not essential to the overall functioning of the

software, but has been chosen in order to provide as robust an implementation as possible. The technique has the added benefit that it can be directed towards any other piece of computer software running on the same system, rather than simply monitoring network traffic. This confers the software with the ability to monitor and control the use of other software tools which might be employed to tamper with its protective functioning, and hence it is able to defend itself against such attack by closing down offending programs and/or logging the tampering attempt.

A preferred embodiment of the software design for a single computer user employing the Windows operating system will now be described, followed by a number of proposed derivative versions intended for use by a range of user groups, with reference to the accompanying drawings in which:

Figure 1 shows a generalised view of typical Internet connection topologies

Figure 2 provides an abstraction of the communication levels used during Internet access

As shown in Figure 1, the connection between a user's computer and the Internet can be divided into four generalised areas, running from the client computer in area A1, via local computer equipment and/or network A2 to a communications link A3, and thence to the Internet Services Provider A4 and out onto the Internet.

Attempts to control Internet access using a proxy server system located at the Internet Services Provider A4 (sometimes referred to as a "walled garden") can possibly be circumvented unless separate dedicated computer hardware is used, and the relevant section of the Internet copied onto it. Such systems are also only able to apply a blanket set of restrictions to all users. Conventional Internet portal sites are located remotely out on the Internet, and so can have no role to play in the restriction of Internet access.

No extant systems attempt to control Internet access by manipulation of the communications link A3 itself, while firewall systems have been used to implement restricted Internet access via the local area network A2. Such systems are only feasible where a client-server network topology is in place and a firewall server or network router installed, and suffer from having been originally designed to screen out intrusion from outside a corporate network, rather than restrict access out onto the Internet from within. Firewall-based systems can, in theory, support multiple user profiles for Internet access, but in practice this often results in performance and scaling problems with the firewall server system, as all network traffic must be monitored with respect to the different user profiles which have been defined.

The software described herein, as with other "personal" Internet filtering systems, works primarily in the client computer area A1, in a range of situations such as home, business and educational use. It would also be feasible for the system's data files and/or executable files to be held on a business server A2 in a client-server environment, or even be located at the Internet Services Provider A4.

As shown in Figure 2, the entire communications channel between the computer user and the Internet can be abstracted into 8 levels, labelled as L1 to L8. The user interface L1 forms the link between the user and any application software L2, which in turn relies on the underlying computer operating system L3. Underpinning the inter-connection of the computer with other systems are a set of globally defined network protocols L4, which are implemented for any specific operating system using software network drivers L5. These allow the computer operating system L3 to communicate with the physical network hardware L6, which connects in turn to a communications link L7 and thence to the Internet Services Provider L8.

Other "personal" Internet filtering systems operate at the network protocol level L4, and carry out their task of monitoring network traffic using techniques such as "packet sniffing", where individual packets of network information are read and analysed for their content. This requires that changes are made to certain key system files on the user's computer, resulting in all such systems being mutually incompatible - they cannot co-exist on the same computer.

The preferred method of operation for the software described here is to monitor activity at the user interface level L1, and then control the interaction between the user and any application software L2 such as a Web browser by also issuing commands at the operating system level L3. This is achieved by using API (Application Programming Interface) programming calls which are provided by the underlying operating system L3 itself as an aid to software testing and debugging. In this example, "call-back functions" are used at level L1 to map the window structure and contents of any running programs, resulting in a data structure which acts as a "fingerprint" by which different application software can be identified and monitored. Once a program has been identified by matching its "fingerprint" with a list of known applications and its operation is being monitored, a number of L3 level programming techniques such as messaging, system hooks and DDE (Dynamic Data Exchange) can be used to exercise control over the operation of the application software. The same technique can also be employed to respond to the use of other software tools which might be used to attempt to tamper with the software installation. Any software development system which allows full access to the underlying operating system API calls can be used to implement the software which has been described, the choice of actual programming language is immaterial.

The initial installation of the software should be arranged so that the computer system is restarted and the software immediately makes one or more hidden copies of itself as an aid to the prevention of tampering by end-users. The software should then be invoked in a "properties" mode where the supervisor is able to carry out functions such as changing the default password, editing system parameters, viewing log files and/or modifying the list of permitted Internet sites and range of categories. Categories and sites may be added, changed and deleted, and the resulting master data files should be stored in encrypted form to avoid tampering. No further protection of the data files is required, as their removal or damage will simply result in a further restriction of the end-user's Internet access.

Once the master set of permitted sites has been defined, the end-user's current set of access permissions can be specified by switching access to any given category or site "on" or "off" using the properties mode of the software described in the previous paragraph. This individual set of permissions is then saved as a separate, encrypted data file, following which the system needs to generate the necessary HTML (Hyper-Text Mark-up Language) code to form the personalised, user-specific Internet portal that the permissions have explicitly defined. Normally, this HTML page or pages will be set up as the starting page for the end-user's Internet browsing software.

Once the end-user's computer has been set up by the supervisor, as described, the browser lock software should be executed automatically whenever the computer is started up or restarted. The software then reads and decrypts the user's permissions file into an internal data structure, along with the necessary "fingerprints" for identifying the applications software which is required to be monitored. Following this, the software runs in the background, waiting for applications software in which it is interested to be started, unless it is deactivated via password-protected supervisor access. When a web browser or similar piece of software is detected, its operation is monitored and attempts to access Internet URLs (Uniform Resource Locators) are compared with the user's list of permitted URLs. It is recommended that this URL matching is implemented on a partial basis, so that any Internet page which is "below" a permitted URL in an Internet site's structure is also automatically permitted. Where a URL does not match with the permissions list, system programming calls are used to prevent users' interface operations (mouse movements and clicks, key presses) reaching the application software, or to instruct the application to terminate loading of the page if it has begun, or to shut down the offending software completely, or any combination of these measures. The option also exists to log a user's Internet access, rather than actively control it, or to simply employ the user-specific portal aspect of the software as a starting point for completely unrestricted Internet access.

Based on the example software implementation which has been described, further example systems can be derived which offer options such as:

1. Multiple user data files on the same machine, for cases, such as family use, where a number of users have access to the same computer.
2. The ability to be installed and/or supervised over a typical small peer-to-peer computer network for business or educational use.
3. The ability to use remote, centralised user data files and/or executable program files in a client-server or network computing environment.

These further examples are intended to illustrate the type of specific software systems which might be constructed around the core concept of fusing the functions of an Internet portal with the task of controlling and monitoring Internet access.

CLAIMS

1. A computer software system which combines the function of providing a computing device user with a structured access point to the Internet with the function of restricting access to the Internet by an end-user, based on parameters specified by some person acting in a supervisory capacity.
2. A computer software system as claimed in Claim 1 where Internet access is restricted on the basis of a list of permitted sites, with all other sites being considered as prohibited.
3. A computer software system as claimed in Claim 1 or Claim 2 where the function of monitoring and/or controlling Internet access is achieved at the user interface or operating system level, without the need to access any underlying network protocols.
4. A computer software system as claimed in Claim 3 where the same techniques are employed to resist attempts by end-users to tamper with the system's functioning.
5. A computer software system as claimed in any preceding claim including the ability to support multiple user profiles on the same computing device.
6. A computer software system as claimed in any preceding claim including the ability for multiple copies of the software to be installed and/or supervised over a computer network.
7. A computer software system as claimed in any preceding claim including the ability for the software to be installed and/or operated using remote, centralised user data files and/or executable program files in a client-server or network computing environment.
8. A computer software system substantially as herein described and illustrated in the accompanying drawings.



Application No: GB 9900884.9
Claims searched: 1

Examiner: Leslie Middleton
Date of search: 13 September 2000

Patents Act 1977
Search Report under Section 17

Databases searched:

UK Patent Office collections, including GB, EP, WO & US patent specifications, in:

UK Cl (Ed.R): G4A (AUDB)

Int Cl (Ed.7): G06F 17/30, H04L 029/10

Other: Online: EPODOC, JAPIO, WPI / EPOQUE; Google / Internet

Documents considered to be relevant:

Category	Identity of document and relevant passage	Relevant to claims
X	WO 98/48546 A1 (Appaloosa Interactive Corp) See pp. 1-9	1,2,3 at least
X	EP 0748095 A2 (AT & T Corp) See pp. 1-6	1 at least
X	PC Magazine: The 1997 Utility Guide (Internet Filtering Utilities) (Editor's Choice) http://www.zdnet.com/pcmag/features/utility/filter/ufuec.htm (Cyber Patrol) (Net Nanny)	1,2,3,5,6 at least
X	Comms of Assocn for Computing Machinery, 1996, Vol 39, No 10, pp.87-93: "PICS:Internet Access Controls without Censorship" Resnick P et al	1

X	Document indicating lack of novelty or inventive step	A	Document indicating technological background and/or state of the art.
Y	Document indicating lack of inventive step if combined with one or more other documents of same category.	P	Document published on or after the declared priority date but before the filing date of this invention.
&	Member of the same patent family	E	Patent document published on or after, but with priority date earlier than, the filing date of this application.